WEST

Freeform Search

Database: US Patents Full-Text Database
US Pre-Grant Publication Full-Text Database
JPO Abstracts Database
EPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Term: (hosts with plurality with (virtual adj1 (wan or lan)))

Display: 10 **Documents in Display Format:** KWIC **Starting with Number** 1

Generate: ○ **Hit List** ◉ **Hit Count** ○ **Side by Side** ○ **Image**

Search | Clear | Help | Logout | Interrupt

Main Menu | Show S Numbers | Edit S Numbers | Preferences | Cases

Search History

DATE: Thursday, December 04, 2003 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=USPT; PLUR=YES; OP=ADJ			
L11	L10	1	L11
L10	L9 and L8	1	L10
L9	(hosts with plurality with (virtual adj1 (wan or lan)))	3	L9
L8	(plurality with (virtual adj1 (wan or lan))).ab.	5	L8
L7	(plurality with (virtual adj1 (wan or lan)))	26	L7
L6	(virtual adj1 (wan or lan))	362	L6
L5	L3 and (plurality with hosts)	2	L5
L4	L3 and hosts	5	L4
L3	(plurality with (virtual adj1 servers))	9	L3
L2	L1 and (plurality with (virtual adj1 servers))	1	L2
L1	(virtual adj1 (wans or lans))	362	L1

END OF SEARCH HISTORY

L9: Entry 2 of 3 File: USPT Sep 1, 1998

DOCUMENT-IDENTIFIER: US 5802047 A
TITLE: Inter-LAN connecting device with combination of routing and switching functions

Detailed Description Text (6):
A plurality of hosts A to N are connected to the Ether switch 3a and the physical ports 33 to 40 and the logic ports are divided into groups to make virtual LANs 3b to 3d. Further, receive buffers 32a, 32c, 32e, 32g, 32i, 32k, 32m and 32o and transmit buffers 32b, 32d, 32f, 32h, 32j, 32l, 32n and 32p within the transmit-receive buffer 32 for the physical ports correspond to the physical ports 33 to 40, respectively.

Tabl

WEST

Generate Collection

L7: Entry 26 of 26 File: USPT Mar 11, 1997

DOCUMENT-IDENTIFIER: US 5610920 A
TITLE: Coupling of voice and computer resources over networks

Brief Summary Text (17):
U.S. Pat. No. 5,351,237, issued to Shinohara, et al. entitled "Network System Comprising a Plurality of LANs Connected to an ISDN via a Plurality of Routers, Each Capable of Automatically Creating a Table for Storing Router Information" This patent describes a bridging system to create a virtual LAN topology between remotely located physical LANs using an integrated services digital network (ISDN) as the network to interconnect the physically disparate LANs. The system describes Internet protocol (IP) address assignments and methods for a routing protocol.

First Hit Fwd Refs

Generate Collection

L14: Entry 4 of 12 File: USPT Jul 31, 2001

DOCUMENT-IDENTIFIER: US 6269404 B1
TITLE: Virtual network architecture for connectionless LAN backbone

Abstract Text (1):
Network traffic management is achieved based on automatically setting up a plurality of virtual networks (VNETs) within a single large virtual LAN. Multicast/broadcast traffic is confined to the VNET of the source, without imposing constraints on layer two addressing within the virtual LAN. VNETs are domains of users of a virtual LAN which include members of logical networks defined at layer three or higher. One method includes transferring a multi-destination packet originating from a particular node in the virtual LAN by tunnelling across a connectionless backbone network to a virtual net server. The virtual net server translates the multi-destination packet to a plurality of tunneled messages identifying nodes authorized to receive multi-destination packets from members of the particular VNET which originated the packet. The tunneled messages are then forwarded from the virtual net server to the authorized nodes. This way, multi-destination packets, such as advertisement or discovery packets, are confined to a single VNET. By confining the multi-destination packets to a single VNET, unicast packets generated within the virtual LAN are then also naturally confined to that VNET.

Brief Summary Text (3):
The present invention relates to data communication networks which consist of a number of local area network (LAN) segments interconnected to form a virtual LAN environment; and more particularly to methods for managing data flow in such networks across a connectionless LAN backbone.

Brief Summary Text (5):
Historically, networks have been designed around the wired LAN segment as the basic technique for establishing network user groups. Standard network layer protocols define logical networks with a single layer two (data link layer) LAN segment in mind, with layer two bridging and layer three (network layer) routing functions used for moving data between LAN segments and layer three logical networks. However, with the emerging ATM LAN emulation mode and other LAN switching systems, the layer two boundaries become less controlled, giving rise to the concept of a virtual LAN. See, U.S. Pat. No. 4,823,338 to Chan et al., and an IEEE standard referred to as 802.1D. Nodes in a single layer two virtual LAN are found on different physical LAN segments but have the appearance to layer two processes (data link layer processes using medium access control MAC addresses) of residing on a single layer two LAN segment. This allows a unicast packet to propagate across the virtual LAN to any other station in the virtual LAN. Also, multi-destination packets generated on a particular LAN segment propagate throughout a number of interconnected LAN segments to ensure that all possible members of the virtual LAN receive the packet.

Brief Summary Text (6):
Within virtual LAN domains, multicast/broadcast frames are used by higher layer "discovery" or "advertisement" procedures to locate other systems or services within the virtual LAN domain. Systems send "data" to other systems using unicast MAC address which are either known in advance or learned through

multicast/broadcast discovery and advertisement procedures. Systems send "multi-media data" using either unicast or multicast frames with special protocols to improve throughput or latency, as required.

Brief Summary Text (7):
Large virtual LANs create large multicast/broadcast domains; and the burden on the backbone network of transmitting all these multi-destination packets begins to impact overall system performance. More importantly, the users of the virtual LAN become burdened by a large number of multi-destination packets that must be inspected and processed, even when the packet is simply discarded. In fact, several layer three network protocols may co-exist in a single virtual LAN, resulting in much traffic which is irrelevant to many users in the virtual LAN, which must nonetheless process the traffic to discover that the network layer data unit carried in it relates to a protocol it does not use.

Brief Summary Text (8):
Commonly used network layer protocols include the internet protocol (IP) originally developed under DARPA, the interpacket exchange protocol (IPX) published by Novell, the Xerox network system (XNS) published by Xerox, the Banyan VINES protocol, the NetBIOS protocol published by IBM and Microsoft, Apple Talk published by Apple Computer, and the DECNet protocol published by Digital Equipment Corporation. Many network layer protocols create protocol specific domains based on the logical network identifiers. For example, the IP protocol establishes "subnet" domains based on the network number portion, and extensions, of the IP address of the frame. The IPX protocol creates logical networks based on the internal network number assigned to servers in the network. Apple Talk creates "zones". The NetBIOS protocol does not support multiple domains within a single LAN or emulated LAN, and can thus be considered to define a single (or "null") logical network at layer three, by default. These protocol specific logical networks defined at layer three, or higher layers, are called virtual networks, or VNETs in the present application. By the nature of virtual LANs according the prior art, the broadcast/multicast boundaries of the virtual LAN and of the VNETs are equal. Thus, as mentioned above, multicast/broadcast traffic for IPX networks will be received and processed by nodes which are members of an IP subnet, if both nodes fall in the same virtual LAN.

Brief Summary Text (9):
Prior art techniques have arisen to divide networks into several virtual LANs. U.S. Pat. No. 5,394,402 to Ross describes a virtual LAN architecture in a network which includes a backbone using a synchronous transfer mode (ATM) switching. The virtual LAN groupings act to limit the size of the multicast/broadcast domains by constraining the layer two addressing within the virtual LAN, and thus help manage the amount of multicast/broadcast packets which must be handled by a user of the network. To cross virtual LAN boundaries, internetworking devices providing layer three routing functions are required. Thus, when a change is made in a network having a number of virtual LANs, such as a new node being added, or a user moving from one LAN segment to another LAN segment in a different virtual LAN, the VNETs must be reconfigured for the new or moved node, such as by assigning a new layer three address to the node and the like. This complication has effects throughout the network, as the internetworking devices in the system need to learn the new information, and to learn that the old information in the case of a moved node, is obsolete. Further, individual users of the virtual LANs which may have cached the old layer two MAC address of the moved node, will lose track of the node, as it will not be able to send a packet across the virtual LAN boundary with the cached layer two MAC address. Also, the use of several virtual LANs within an organization, may place constraints on layer three network definition. For instance, the IPX network number used in the VNET of a first virtual LAN should not be used in the VNET of a second virtual LAN, because if a node moves from the first to the second, the moved node might erroneously access resources in the VNET of new virtual LAN with the network number of the VNET in old virtual LAN.

h e b b g e e e f e e e g

Brief Summary Text (12):
According to the present invention, network traffic management is achieved based on automatically setting up a plurality of VNETs within a single large virtual LAN. Multicast/broadcast traffic is confined to the VNET of the source, without imposing constraints on layer two addressing within the virtual LAN. Thus, when a node is moved within the network from one segment to another, it remains within the same virtual LAN, so that it may keep its layer three address or addresses, and unicast packets addressed to it from other users of the virtual LAN find their destination. Furthermore layer three network configuration is unconstrained.

Brief Summary Text (13):
The present invention can be characterized as a method for managing traffic in a network based on a set of local area network segments interconnected as a virtual LAN, and in which nodes on respective LAN segments in the set are members of VNETs. The method includes tunneling a multi-destination packet originating from a particular node in the virtual LAN, encapsulated, or otherwise reformatted, as a single destination message to a virtual net server. The virtual net server translates the multi-destination packet to a plurality of directed messages identifying nodes authorized to receive multi-destination packets from members of the particular VNET which originated the packet. The directed messages are then forwarded from the virtual net server to the authorized nodes. This way, multi-destination packets, such as advertisement or discovery packets, are confined to a single VNET. By confining the multi-destination packets to a single VNET, unicast packets generated within the virtual LAN are then also naturally confined to that VNET. Packets are naturally confined to the VNET, because the advertisement of their address, and the procedures used to discover the addresses of others, are prevented from exiting the-VNET of the particular node which issues the multi-destination packet. The present invention elegantly controls proliferation of multicast/broadcast traffic in large virtual LANs and confines unicast traffic to the VNET of the source, without introducing the complexities of prior art techniques to divide large virtual LANs into several smaller ones.

Brief Summary Text (14):
According to this aspect of the invention, the virtual net server automatically configures itself in response to the multi-destination packets received at the virtual net server, and in response to the layer three networks set up in the virtual LAN. Thus, when a virtual net server receives a multi-destination packet, it determines a virtual net domain based on the layer three network protocol and logical network which originated the packet, and the source medium access control (MAC) address of the packet. If a packet is received from a source node which had not previously sent a packet using the identified logical network identifier, then a connection is set up between the virtual net server and that source node, adding the new node to the appropriate virtual net domain. Thus, the virtual net domain is defined as a group of nodes intended to receive multi-destination packets from members of a particular VNET determined by a layer three network protocol/network identifier.

Brief Summary Text (24):
Accordingly, a low-cost virtual LAN/virtual net (VLAN/VNET) architecture has been provided. Edge devices operate at layer 2, based on MAC address filtering. The layer 3 multi-protocol complexities are confined in the virtual net server on the backbone LAN. However, the layer 3 multi-protocol complexity only includes components necessary to decode and forward the multi-destination frames. Furthermore, the virtual NET server, the edge devices and adapters automatically learn virtual net domains of LAN segments and nodes in the system.

Brief Summary Text (26):
The present invention greatly improves flexibility of network architectures by managing the flow of traffic within virtual LANs. The invention allows the creation

of a plurality of VNETs within the virtual LAN according to guidelines unique to each installation, such as shared access to services, using existing logical network constructs of standard layer three protocols.

Drawing Description Text (2):
FIG. 1 provides a conceptual overview of a network configured with virtual LAN domains and virtual net domains according to the present invention.

Drawing Description Text (6):
FIG. 5 is a flow chart of the basic tunnelling process used in the system of FIGS. 3 and 4.

Drawing Description Text (7):
FIGS. 6 and 7 illustrate encapsulation of multicast packets for the tunnelling process.

Detailed Description Text (2):
FIG. 1 provides a conceptual overview of a network in which the present invention operates. The network includes a plurality of LAN segments coupled to end systems or nodes on the network. The LAN segments include segments 10-17 which are connected to an edge device 18, and segments 19-26 which are coupled to edge device 27. A backbone network 28 is coupled to each of the edge devices 18 and 27 to provide interconnection among the LAN segments. Also coupled to the backbone network, may be adapters, such as adapters 30 and 31 which connect directly to end systems. Within the wired network which includes the LAN segments, edge devices, adapters and the backbone, a virtual LAN domain 35 may be established as a layer two construct. Many virtual LANs may be implemented using higher layer procedures, such as described in the Ross patent referred to above or otherwise, but the example of one virtual LAN is used to illustrate the present invention.

Detailed Description Text (3):
According to the present invention, multiple virtual net domains, including virtual net domain C, virtual net domain B and virtual net domain A are set up within a single virtual LAN domain 35. A virtual net domain is defined as the set of LAN segments/ATM systems that are members of the same network layer protocol logical networks which are identified by a unique network layer identifier, and may be extended to include other nodes intended to receive packets from members of this logical network.

Detailed Description Text (4):
Virtual LAN domains contain numerous interconnected LAN segments, each with one or more attached systems (desktops, servers, routers, etc.) interconnected across a backbone 28. The utilized protocol stacks within the network (e.g. IP, IPX) must be able to function properly within the virtual LAN domain.

Detailed Description Text (5):
A virtual network configuration is utilized when variant network layer protocols, and logical networks are used within the virtual LAN. For example, a single virtual LAN wide virtual net domain may be created for IP, while requiring creation of several IPX virtual net domains. Each LAN segment end system can be then individually attached to differing IPX VNETs based on policies such as desired services. Also, separate VNET domains may be created for many IP subnets and many IPX networks. Each LAN segment end system can be then individually attached to both IP and IPX VNET domains based on policies such as desired services. Within a single virtual LAN domain, LAN segments and end systems may attach to multiple VNET domains.

Detailed Description Text (6):
According to the present invention, the flow of multicast /broadcast MAC frames are kept within the associated VNET domain. All unicast MAC frames are sent across the

standard virtual LAN. However, the interesting point is that each unicast frame will be addressed at layer 2 to stay within its VNET domain automatically. This means the edge devices are fast, inexpensive and simple while the VLAN/VNET architecture is optimized.

Detailed Description Text (10):
As represented by the cloud 105, a variety of other edge devices and ATM adapters may be coupled to the LAN emulation backbone 100, to establish a virtual LAN over a wide variety of LAN segments and across wide area links.

Detailed Description Text (12):
FIG. 3 illustrates the architecture of a network using virtual net architecture according to the present invention with a backbone network which may be implemented using a connectionless protocol such as FDDI, Ethernet or Token Ring. Thus, as can be seen in FIG. 3, a backbone network 120 is coupled to a first edge device 121 and a second edge device 122. The edge device 121 includes ports P1-PN as illustrated in the Fig. Each of the ports is coupled to a corresponding LAN segment executing a connectionless protocol such as specified in 802.X standards or other protocols. Alternatively, one or more ports may be connected to an ATM edge device which extends the virtual LAN across an ATM emulation backbone.

Detailed Description Text (13):
In the same manner, edge device 122 has a plurality of ports P1-PN as shown in FIG. 3, coupled to a variety of LAN segments. Also, an end system 123 may be coupled directly to the backbone network 120. As represented by the cloud 124, a variety of other edge devices and end systems may be coupled to the backbone 120, to establish a virtual LAN over a wide variety of LAN segments and across wide area links.

Detailed Description Text (14):
According to the present invention, a virtual LAN and virtual net server 125 (VLAN.backslash.VNET server) is coupled to the backbone network 120, such as in an end system on the backbone network, or in a network intermediate system device like a router, bridge or switch on the backbone network 120. Also, virtual net agents 127 and 128 are implemented in the edge devices 121 and 122 respectively. When a multicast frame is detected on a LAN segment in an edge device 121, the multi-cast packet is tunneled to the VLAN.backslash.VNET server 125 through the backbone network. The server 125 translates the multi-destination packet into a plurality of tunneled messages which are sent to virtual net agents 127 and 128 in the edge devices coupled to the backbone network 120. The virtual net agents 127 and 128 then forward the multi-destination packet out port of the edge device on which nodes authorized to receive the multi-destination packet are found. When there is one user on each LAN segment, the multi-destination packet can be delivered exclusively to members of the virtual net domain using this architecture.

Detailed Description Text (16):
The server 201 includes a decoder 203, and a plurality of virtual net tunnel modules 204, 205, 206, 207. The edge device 200 includes an agent 208 which operates with the server 201.

Detailed Description Text (17):
In operation, an edge device 200 receives on an incoming port 210 from an originating LAN segment a multicast or broadcast packet. This packet is then routed using tunneling 211 under control of the agent 208 to the server 201.

Detailed Description Text (18):
The decoder 203 in the virtual net/virtual LAN server 201 (V/V server) determines the virtual net domain of the packet in response to the network protocol of the multicast/broadcast packet and the network identifier, if used in the identified protocol, by layer three protocol constructs in the packet. It then passes the multi-destination packet to the appropriate virtual net tunnel module. A virtual

net domain exists for each network address value (for example each IP subnet value) supported by a given network protocol. When the frame does not contain a network identifier (for instance a NetBIOS frame) only one virtual net tunnel module exists for it in the server 201. Thus, if the multicast/broadcast packet is an IPX packet, then it is forwarded across either line 212 or line 213 to the VNET tunnel module 206 or module 207 for corresponding network identifiers. If the multicast/broadcast packet is an IP protocol packet, then it is forwarded across line 214 or line 215 to one of the VNET tunnel modules 204 or 205 for corresponding subnets. For the purposes of this example, the packet is passed to the subnet 1 tunnel module 204. The subnet 1 tunnel module 204 includes a table 216 which maps the virtual net domain to established tunnels in the backbone. Established tunnels provide direct paths (e.g. 217) by means of single destination packets to agents in edge devices 200 on which ports authorized to receive the frame are found. Thus, in this example, the agent 208 in the edge device 200 receives the multicast packet across tunnel 217 and forwards the packet out the appropriate ports. The agent includes table 218 which maps the tunnel on which the multicast packet is received, to ports on which nodes authorized to receive the packet are found, using the source MAC address of the packet to make sure that it is not sent back on the originating LAN segment 210. In the illustrated example, the packet is sent on port 219 and port 220 by the agent 208, but not on other ports of the edge device and not on the port coupled to originating LAN segment 210.

Detailed Description Text (19):
Alternatively, a VNET encapsulation can be utilized which provides the ability for each of the virtual net tunnel modules to share tunnel addresses established for each of the edge devices. However, the agent at the edge device must be able to handle the decapsulation of the frame as it received and route it appropriately.

Detailed Description Text (20):
Standard end system adapters can utilize this virtual net architecture. The configuration steps are not required for end systems because there is no edge device connected to them which operates as a proxy for other LAN segments. Since the directly connected end systems do not register as proxies within the server, the end system is automatically distinguished from an edge device. Thus, only one tunnel per virtual LAN is set up from the server to the end system. All associated virtual net membership entries set the tunnel identifier values to the same tunnel for the adapter. In the configuration process, the virtual net membership entry, layer two MAC address is not forwarded to the end system adapter, because it is not necessary for use there.

Detailed Description Text (22):
The basic tunneling process according to the present invention is illustrated with respect to FIGS. 5, 6 and 7. In FIG. 5, the basic process is described, which begins with receiving an incoming multicast packet at a virtual net agent on an edge device. The edge device forwards the multicast packet on attached segments, with or without filtering by virtual net domain (block 150). Alternatively, it may defer forwarding the packet on attached segments until it receives the multicast back from the server.

Detailed Description Text (23):
Next, the virtual net agent encapsulates the multicast packet and tunnels it to the virtual net/virtual LAN server on the backbone network (block 151). At the virtual LAN/virtual net server, the multicast packet is decapsulated from the tunneled message, and encapsulated in a new tunnel packet for forwarding to other virtual net agents (block 151). At the agents, the tunneled messages are decapsulated, and the outgoing multicast packets is forwarded on attached segments, other than the segment which originated the message and segments which already received the message (block 153).

Detailed Description Text (24):

h e b b g e e e f e e e g

The tunneling process can be understood with references to FIGS. 6 and 7. Basically, tunneling involves encapsulating a multi-destination packet in a single destination packet having a MAC address of the destination of the tunnel, and a source address equal to the source of the tunnel. Thus, the agent in the edge device will encapsulate the message as illustrated in FIG. 6 where the multicast frame 155 is encapsulated in a single destination packet having the server address 156 as a destination address and other supporting control fields, such as the frame check sequence 157, surrounding the multicast frame 155. At the receiving end of the tunnel, the server receives the frame and processes it. It discovers that the frame is a tunneled multicast frame, and using the process described with respect to FIG. 4, encapsulates the frame in a tunnel directed to the agent as shown in FIG. 7. Thus, a tunnel from the server to the agent will carry destination address equal to the agent address 160, the multicast frame 155 will be encapsulated within the packet. Supporting control fields, such as the frame check sequence 161 and the like, are included within the tunneled packet.

Detailed Description Text (25):
The tunneling process can take a variety of formats. For instance, each tunnel may be established by setting up a specific destination and source address for each tunnel handler/edge device pair. This way, the agent and server must maintain a number source and destination addresses, and correlate those with specific VNETs. Alternatively, a single address may be used for the server and a single address used for each agent, and the tunnel packet will carry control fields which specify the information needed to recognize the packet as a tunneled packet.

Detailed Description Text (26):
FIG. 8 provides a flow chart illustrating the handling of multicast and broadcast packets according to a centralized server embodiment of this system. The process begins when an edge device receives a frame (block 250). The device determines the type of frame (block 251). If frame is a unicast frame, then it is handled with standard LAN techniques (block 252). If the frame is a multicast or broadcast, then the edge device forwards the frame using tunneling to the V/V server (block 254).

Detailed Description Text (27):
In the server, the decoder determines the virtual net domain of the frame (block 255). Then the frame is passed to the determined virtual net tunnel handler (block 256). The virtual net tunnel handler determines whether the source of the multicast packet has a corresponding entry in its membership list (block 257). If an entry is found, then the frame is forwarded as a tunneled message as set up in the membership list (block 258).

Detailed Description Text (28):
If no entry was found in the membership list at the virtual net tunnel handler in the test of block 257, then an automatic configuration routine is executed (block 260). After the configuration routine, then the process proceeds to block 258 to forward the frame across the established tunnels for the virtual net tunnel handler.

Detailed Description Text (29):
The process of block 258 results in edge devices receiving the multicast/broadcast frame. Each edge device which receives the multicast/broadcast frame, then sends the frame once on ports to members of the virtual net domain. This is done by the edge device maintaining a table which maps the tunnel on which the frame is received to specific ports, or modules accessible through the ports, of the edge device. However, the edge device does not send the multicast packet back out on the segment which originated the packet. This is determined by checking the source address of the multicast/broadcast frame, and comparing that source address with the address of devices on the respective ports (block 259).

Detailed Description Text (30):

h e b b g e e e f e e e g

FIG. 9 illustrates the configuration routine executed at block 260 of FIG. 8. According to this routine, the virtual net tunnel handler sets up a tunnel from the virtual net tunnel handler to the originating edge device (block 270). An entry is created in the virtual net membership list for the source of the packet (block 271). The entry includes a source MAC address of the originating end station and a tunnel identifier (i.e. MAC address of the agent in the edge device) (block 272). After creating the entry, the virtual net tunnel handler sends the source MAC address across the established tunnel to the originating edge device (block 273). The edge device then stores the received source MAC address in the virtual channel/virtual net membership list maintained by the agent (block 274). This MAC address is utilized to map incoming frames on this tunnel to the appropriate ports of the edge device.

Detailed Description Text (31):
Utilizing the process of FIGS. 8 and 9, each virtual net tunnel handler 204, 205, 206, 207 as shown in FIG. 4, establishes a tunnel to each edge device which includes a port through which a member of the virtual net domain is found. These established tunnels provide a mechanism for distributing the multiple destination packets efficiently across the backbone. The decoder in the server 201 maps the incoming packet to the virtual net tunnel handler which maps the frame based on a membership list to a set of established tunnels. The edge devices map frames incoming on specific tunnels to ports of the edge device. This tightly controls the propagation of multiple destination packets within the appropriate virtual net domain of the originating device.

Detailed Description Text (32):
As mentioned above, the V/V server can be distributed to the edge devices, rather than executed in a centralized site. FIGS. 10 and 11 illustrate a process which is executed in the edge devices according to this distributed virtual net server model. Thus, in FIG. 10, the process executed by the edge device when it receives a packet from the LAN segment on the user side of the edge device is shown. The process begins with receiving a packet from the user side (block 300). The process then determines the type of frame (block 301). If it is a unicast frame, then it is handled with standard LAN procedures (block 302). If it is a multicast or broadcast packet, the distributed virtual net server determines the VNET domain of the frame (block 303). The source address of the multicast packet is added to a VNET domain list which is maintained in the edge device, if it is not already there (block 304). Finally, a multicast packet is forwarded using tunneling to other edge devices, and sent on local LAN segments. Optionally, transmission on the local LAN segments can be filtered by VNET membership (block 305).

Detailed Description Text (33):
The process shown in FIG. 11 is executed by edge devices receiving a packet from a tunnel from other edge devices. Thus, the process begins with receiving a frame from the tunnel on the backbone network (block 310). The edge device next determines the type of frame (block 311). If it is a unicast frame, it is handled with the standard LAN processes (block 312). If the frame is a multicast or a broadcast packet, the edge device determines whether the packet has a local source address (block 313). If it has a local source address, then it is discarded (block 314), because the process described in FIG. 10 has already forwarded the packet to the local LAN segments. If the packet does not have a local source address, then the edge device sends the frame once on ports coupled to members of the virtual net membership list which matches the multicast packet (block 315). Accordingly, the virtual net server can be distributed to the edge devices in the virtual LAN architecture.

Detailed Description Text (34):
The present invention provides management of traffic in a virtual LAN environment according to the concept of a virtual net domain. To maintain virtual net domain boundaries, edge devices operate at layer 2, while limited layer 3 complexity is

centralized in, for example, an improved server, or alternatively in a distributed virtual net server. The virtual net server, the edge devices and adapters automatically learn virtual net domain membership among nodes on connected LAN segments.

Current US Original Classification (1):
709/238

Current US Cross Reference Classification (1):
709/245

Other Reference Publication (4):
Simpson, W., "IP in IP Tunneling", Request for Comments #1853, Oct. 1995, 9 pages.

CLAIMS:

1. For a network including a set of local area network (LAN) segments interconnected as a virtual LAN in which nodes in the virtual LAN are members of one or more logical networks, a method for managing traffic in the network, comprising:

detecting a broadcast packet on a LAN segment within the set;

supplying the detected broadcast packet to a server;

determining the logical network for which the detected broadcast packet is intended based upon a comparison of a detected broadcast packet's source against a membership list; and

confining the broadcast packet to nodes in the determined logical network.

2. The method of claim 1, including after said step of detecting:

tunneling the detected broadcast packet by a single destination packet to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the broadcast packet.

3. The method of claim 1, wherein the network includes a connectionless backbone communication path and a plurality of edge devices which interconnect the set of LAN segments and the backbone communication path, and including:

detecting in a particular edge device, a broadcast packet on a particular LAN segment in response to a medium access control MAC layer address in the broadcast packet;

tunneling the broadcast packet detected in the particular edge device from the particular edge device across the backbone communication path to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the broadcast packet detected in the particular edge device.

7. A method of communication comprising:

interconnecting a set of local area network (LAN) segments as a virtual LAN in which nodes in the virtual LAN are members of one or more logical networks, the logical networks defined at a particular layer, the particular layer being different from a layer at which the virtual LAN is defined;

h e b b g e e e f e e e g

detecting a multicast packet on a LAN segment within the set, the multicast packet having a multi-destination address of a layer lower than the particular layer and defining a set of recipients;

supplying the detected multicast packet to a server;

determining the logical network for which the detected multicast packet is intended based upon a comparison of a detected multicast packet's source against a membership list, the logical network including a subset of the set of recipients; and

confining the multicast packet to nodes in the virtual LAN authorized to receive multi-destination packets intended for members of the determined logical network.

10. The method of claim 7, including after said step of detecting:

tunneling the detected multicast packet by a single destination packet to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the multicast packet.

11. The method of claim 7, wherein the Virtual LAN includes a connectionless backbone communication path and a plurality of edge devices which interconnect the set of LAN segments and the backbone communication path, and including:

detecting in a particular edge device, a multicast packet on a particular LAN segment in response to a medium access control MAC layer address in the multicast packet;

tunneling the multicast packet detected in the particular edge device from the particular edge device across the backbone communication path to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the multicast packet detected in the particular edge device.

15. For a network including a set of local area network (LAN) segments interconnected as a virtual LAN in which nodes in the virtual LAN are members of one or more logical networks, a method for managing traffic in the network, comprising:

detecting a broadcast packet on a LAN segment within the set, the broadcast packet having a MAC layer broadcast address;

supplying the detected broadcast packet to a server;

determining the logical network for which the detected broadcast packet is intended based upon a comparison of a detected broadcast packet's source against a membership list; and

confining the broadcast packet to nodes in the virtual LAN authorized to receive multi-destination packets intended for members of the determined logical network.

16. The method of claim 15, including after said step of detecting:

tunneling the detected broadcast packet by a single destination packet to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the broadcast packet.

17. The method of claim 15, wherein the network includes a connectionless backbone communication and a plurality of edge devices which interconnect the set of LAN segments and the backbone communication path, and including:

detecting in a particular edge device, a broadcast packet on a particular LAN segment in response to a medium access control MAC layer address in the broadcast packet;

tunneling the broadcast packet detected in the particular edge device from the particular edge device across the backbone communication path to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the broadcast packet detected in the particular edge device.

21. A method of communication comprising:

interconnecting a set of local area network (LAN) segments as a virtual LAN in which nodes in the virtual LAN are members of one or more logical networks, the logical networks defined in layer three or higher;

detecting a multicast packet on a LAN segment within the set, the multicast packet having a MAC multi-destination address defining a set of recipients;

supplying the detected multicast packet to a server;

determining the logical network for which the detected multicast packet is intended based upon a comparison of a detected multicast packet's source against a membership list, the logical network including a subset of the set of recipients; and

confining the multicast packet to nodes in the determined logical network.

22. The method of claim 21, including after said step of detecting:

tunneling the detected multicast packet by a single destination packet to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the multicast packet.

23. The method of claim 21, wherein the network includes a connectionless backbone communication path and a plurality of edge devices which interconnect the set of LAN segments and the backbone communication path, and including:

detecting in a particular edge device, a multicast packet on a particular LAN segment in response to a medium access control MAC layer address in the multicast packet;

tunneling the multicast packet detected in the particular edge device from the particular edge device across the backbone communication path to the server; and

producing a plurality of single destination messages in the server to deliver the information carried by the multicast packet detected in the particular edge device.

First Hit Fwd Refs

Generate Collection

L5: Entry 7 of 12 File: USPT Apr 4, 2000

DOCUMENT-IDENTIFIER: US 6047325 A
TITLE: Network device for supporting construction of virtual local area networks on arbitrary local and wide area computer networks

Abstract Text (1):
A network device that translates addresses of machines on physically separate networks and filters packets at the link, network and transport layers implements a virtual LAN over interconnected computer networks transparent to the computer networks. Using authentication and encryption, a secure connection between these network devices over a public wide area network implements a virtual private network and enables the definition of virtual LANs over the virtual private network. The network device has three tables for network address translation, routing, and filtering. A controller processes each incoming packet by translating network addresses to determine the destination of the packet, routing the packet to the determined location and filtering the packet according to filters defined for traffic between the source destination of the packet. If the packet is to be directed to a wide area network, encryption and authentication procedures can be provided to ensure secure transmission of the packet.

Current US Original Classification (1):
709/227

Current US Cross Reference Classification (1):
709/238

Current US Cross Reference Classification (2):
709/239

Current US Cross Reference Classification (3):
709/245

First Hit Fwd Refs

Generate Collection

L5: Entry 11 of 12 File: USPT Jan 18, 2000

DOCUMENT-IDENTIFIER: US 6016318 A
TITLE: Virtual private network system over public mobile data network and virtual LAN

Abstract Text (1):
In a virtual private network system accessed by an internet, a virtual local area network (LAN) is connected to a LAN emulation server and IAN emulation clients, and a router is connected between the internet and the virtual LAN. Also, a public mobile data network is connected to a location register and mobile data subscriber processing units, and a data gateway is connected between the internet and the public mobile data networks Further, a virtual private network gateway is connected between the virtual LAN and the public mobile data network. A mobile data terminal having one IP address and one public network address and can be connected to either one of the LAN emulation clients or one of the mobile data subscriber processing units.

Current US Cross Reference Classification (3):
709/249

First Hit Fwd Refs

Generate Collection

L14: Entry 12 of 12 File: USPT May 12, 1998

DOCUMENT-IDENTIFIER: US 5752003 A
TITLE: Architecture for managing traffic in a virtual LAN environment

Abstract Text (1):
Network traffic management is achieved based on automatically setting up a plurality of virtual networks (VNETs) within a single large virtual LAN. Multicast/broadcast traffic is confined to the VNET of the source, without imposing constraints on layer two addressing within the virtual LAN. VNETs are domains of users of a virtual LAN which include members of logical networks defined at layer three or higher. One method includes transferring a multi-destination packet originating from a particular node in the virtual LAN by a point-to-point path to a virtual net server. The virtual net server translates the multi-destination packet to a plurality of directed messages identifying nodes authorized to receive multi-destination packets from members of the particular VNET which originated the packet. The directed messages are then forwarded from the virtual net server to the authorized nodes. This way, multi-destination packets, such as advertisement or discovery packets, are confined to a single VNET. By confining the multi-destination packets to a single VNET, unicast packets generated within the virtual LAN are then also naturally confmed to that VNET.

Brief Summary Text (3):
The present invention relates to data communication networks which consist of a number of local area network (LAN) segments interconnected to form a virtual LAN environment; and more particularly to methods for managing data flow in such networks.

Brief Summary Text (5):
Historically, networks have been designed around the wired LAN segment as the basic technique for establishing network user groups. Standard network layer protocols define logical networks with a single layer two (data link layer) LAN segment in mind, with layer two bridging and layer three (network layer) routing functions used for moving data between LAN segments and layer three logical networks. However, with the emerging ATM LAN emulation mode and other LAN switching systems, the layer two boundaries become less controlled, giving rise to the concept of a virtual LAN. See, U.S. Pat. No. 4,823,338 to Chan et al., and an IEEE standard referred to as 802.1D. Nodes in a single layer two virtual LAN are found on different physical LAN segments but have the appearance to layer two processes (data link layer processes using medium access control MAC addresses) of residing on a single layer two LAN segment. This allows a unicast packet to propagate across the virtual LAN to any other station in the virtual LAN. Also, multi-destination packets generated on a particular LAN segment propagate throughout a number of interconnected LAN segments to ensure that all possible members of the virtual LAN receive the packet.

Brief Summary Text (6):
Within virtual LAN domains, multicast/broadcast frames are used by higher layer "discovery" or "advertisement" procedures to locate other systems or services within the virtual LAN domain. Systems send "data" to other systems using unicast MAC address which are either known in advance or learned through multicast /broadcast discovery and advertisement procedures. Systems send "multi-

media data" using either unicast or multicast frames with special protocols to improve throughput or latency, as required.

Brief Summary Text (9):
Large virtual LANs create large multicast/broadcast domains; and the burden on the backbone network of transmitting all these multi-destination packets begins to impact overall system performance. More importantly, the users of the virtual LAN become burdened by a large number of multi-destination packets that must be inspected and processed, even when the packet is simply discarded. In fact, several layer three network protocols may co-exist in a single virtual LAN, resulting in much traffic which is irrelevant to many users in the virtual LAN, which must nonetheless process the traffic to discover that the network layer data unit carried in it relates to a protocol it does not use.

Brief Summary Text (10):
Commonly used network layer protocols include the internet protocol (IP) originally developed under DARPA, the interpacket exchange protocol (IPX) published by Novell, the Xerox network system (XNS) published by Xerox, the Banyan VINES protocol, the NetBIOS protocol published by IBM and Microsoft, AppleTalk published by Apple Computer, and the DECNet protocol published by Digital Equipment Corporation. Many network layer protocols create protocol specific domains based on the logical network identifiers. For example, the IP protocol establishes "subnet" domains based on the network number portion, and extensions, of the IP address of the frame. The IPX protocol creates logical networks based on the internal network number assigned to servers in the network. AppleTalk creates "zones". The NetBIOS protocol does not support multiple domains within a single LAN or emulated LAN, and can thus be considered to define a single (or "null") logical network at layer three, by default. These protocol specific logical networks defined at layer three, or higher layers, are called virtual networks, or VNETs in the present application. By the nature of virtual LANs according the prior art, the broadcast/multicast boundaries of the virtual LAN and of the VNETs are equal. Thus, as mentioned above, multicast/broadcast traffic for IPX networks will be received and processed by nodes which are members of an IP subnet, if both nodes fall in the same virtual LAN.

Brief Summary Text (11):
Prior art techniques have arisen to divide networks into several virtual LANs. U.S. Pat. No. 5,394,402 to Ross describes a virtual LAN architecture in a network which includes a backbone using asynchronous transfer mode (ATM) switching. The virtual LAN groupings act to limit the size of the multicast/broadcast domains by constraining the layer two addressing within the virtual LAN, and thus help manage the amount of multicast/broadcast packets which must be handled by a user of the network.

Brief Summary Text (12):
To cross virtual LAN boundaries, internetworking devices providing layer three routing functions are required. Thus, when a change is made in a network having a number of virtual LANs, such as a new node being added, or a user moving from one LAN segment to another LAN segment in a different virtual LAN, the VNETs must be reconfigured for the new or moved node, such as by assigning a new layer three address to the node and the like. This complication has effects throughout the network, as the intemetworking devices in the system need to learn the new information, and to learn that the old information in the case of a moved node, is obsolete. Further, individual users of the virtual LANs which may have cached the old layer two MAC address of the moved node, will lose track of the node, as it will not be able to send a packet across the virtual LAN boundary with the cached layer two MAC address. Also, the use of several virtual LANs within an organization, may place constraints on layer three network definition. For instance, the IPX network number used in the VNET of a first virtual LAN should not be used in the VNET of a second virtual LAN, because if a node moves from the first

to the second, the moved node might erroneously access resources in the VNET of new virtual LAN with the network number of the VNET in old virtual LAN.

Brief Summary Text (15):
According to the present invention, network traffic management is achieved based on automatically setting up a plurality of VNETs within a single large virtual LAN. Multicast/broadcast traffic is confined to the VNET of the source, without imposing constraints on layer two addressing within the virtual LAN. Thus, when a node is moved within the network from one segment to another, it remains within the same virtual LAN, so that it may keep its layer three address or addresses, and unicast packets addressed to it from other users of the virtual LAN find their destination. Furthermore layer three network configuration is unconstrained.

Brief Summary Text (16):
The present invention can be characterized as a method for managing traffic in a network based on a set of local area network segments interconnected as a virtual LAN, and in which nodes on respective LAN segments in the set are members of VNETS. The method includes transferring a multi-destination packet originating from a particular node in the virtual LAN by a point-to-point path to a virtual net server. The virtual net server translates the multi-destination packet to a plurality of directed messages identifying nodes authorized to receive multi-destination packets from members of the particular VNET which originated the packet. The directed messages are then forwarded from the virtual net server to the authorized nodes. This way, multi-destination packets, such as advertisement or discovery packets, are confined to a single VNET. By confining the multi-destination packets to a single VNET, unicast packets generated within the virtual LAN are then also naturally confined to that VNET. Packets are naturally confined to the VNET, because the advertisement of their address, and the procedures used to discover the addresses of others, are prevented from exiting the VNET of the particular node which issues the multi-destination packet. The present invention elegantly controls proliferation of multicast/broadcast traffic in large virtual LANs and confines unicast traffic to the VNET of the source, without introducing the complexities of prior art techniques to divide large virtual LANs into several smaller ones.

Brief Summary Text (17):
According to one aspect of the invention, the virtual net server automatically configures itself in response to the multi-destination packets received at the virtual net server, and in response to the layer three networks set up in the virtual LAN. Thus, when a virtual net server receives a multi-destination packet, it determines a virtual net domain based on the layer three network protocol and logical network which originated the packet, and the source medium access control (MAC) address of the packet. If a packet is received from a source node which had not previously sent a packet using the identified logical network identifier, then a connection is set up between the virtual net server and that source node, adding the new node to the appropriate virtual net domain. Thus, the virtual net domain is defined as a group of nodes intended to receive multi-destination packets from members of a particular VNET determined by a layer three network protocol/network identifier.

Brief Summary Text (25):
The plurality of directed messages are composed by establishing for the particular virtual LAN, one virtual channel connection for each virtual net domain, to each edge device through which nodes that are members of the particular virtual net domain are accessible. The multi-destination packet is segmented, if needed, into a stream of fixed length ATM packets to form a message. The plurality of directed messages are associated with the established virtual connections. In the edge device, the virtual channel connection is mapped to ports of the edge device through which nodes are accessible that are members of the particular virtual net domain associated with the virtual channel. This mapping may be done in response to

the source address of the multi-destination packet during the configuration process used to set up the virtual channel connection as discussed above.

Brief Summary Text (27):
Alternative systems may use other types of backbone networks, including connectionless backbone networks. In a connectionless backbone network, the multi-destination packets which are detected at edge devices, need to be tunneled, or encapsulated in a single-destination packet, to the virtual net server.

Brief Summary Text (28):
Accordingly, a low-cost virtual LAN/virtual net (VLAN/VNET) architecture has been provided, particularly suited to the ATM standard LAN emulation mode environment. Edge devices operate at layer 2, based on MAC address filtering. The layer 3 multi-protocol complexities are confined in the virtual net server, such as in an improved BUS module according the ATM standard. However, the layer 3 multi-protocol complexity only includes components necessary to decode and forward the multi-destination frames. Standard LAN emulation adapters are supported with no changes in this environment. Furthermore, the improved BUS, the edge devices and adapters automatically learn virtual net domains of LAN segments and nodes in the system.

Brief Summary Text (30):
The present invention greatly improves flexibility of network architectures by managing the flow of traffic within virtual LANs. In networks using ATM LAN emulation or other virtual LAN approaches, the invention allows the creation of a plurality of VNETs within the virtual LAN according to guidelines unique to each installation, such as shared access to services, using existing logical network constructs of standard layer three protocols.

Drawing Description Text (2):
FIG. 1 provides a conceptual overview of a network configured with virtual LAN domains and virtual net domains according to the present invention.

Detailed Description Text (2):
FIG. 1 provides a conceptual overview of a network in which the present invention operates. The network includes a plurality of LAN segments coupled to end systems or nodes on the network. The LAN segments include segments 10-17 which are connected to an edge device 18, and segments 19-26 which are coupled to edge device 27. A backbone network 28 is coupled to each of the edge devices 18 and 27 to provide interconnection among the LAN segments. Also coupled to the backbone network, may be adapters, such as ATM adapters 30 and 31 which connect directly to end systems. Within the wired network which includes the LAN segments, edge devices, adapters and the backbone, a virtual LAN domain 35 may be established as a layer two construct. Many virtual LANs may be implemented using higher layer procedures, such as described in the Ross patent referred to above or otherwise, but the example of one virtual LAN is used to illustrate the present invention.

Detailed Description Text (3):
According to the present invention, multiple virtual net domains, including virtual net domain C, virtual net domain B and virtual net domain A are set up within a single virtual LAN domain 35. A virtual net domain is defined as the set of LAN segments/ATM systems that are members of the same network layer protocol logical networks which are identified by a unique network layer identifier, and may be extended to include other nodes intended to receive packets from members of this logical network.

Detailed Description Text (4):
Virtual LAN domains contain numerous interconnected LAN segments, each with one or more attached systems (desktops, servers, routers, etc.) interconnected across a backbone 28. ATM based VLAN domains support numerous interconnected LAN segments, as well as, directly attached ATM systems across an ATM LAN emulation backbone 28.

h e b b g e e e f e e e g

The utilized protocol stacks within the network (e.g. IP, IPx) must be able to function properly within the virtual LAN domain.

Detailed Description Text (5):
A virtual network configuration is utilized when variant network layer protocols, and logical networks are used within the virtual LAN. For example, a single virtual LAN wide virtual net domain may be created for IP, while requiring creation of several IPX virtual net domains. Each LAN segment/ATM system can be then individually attached to differing IPX VNETs based on policies such as desired services. Also, separate VNET domains may be created for many IP subnets and many IPX networks. Each LAN segment/ATM system can be then individually attached to both IP and IPX VNET domains based on policies such as desired services. Within a single virtual LAN domain, LAN segments and ATM systems may attach to multiple VNET domains.

Detailed Description Text (6):
According to the present invention, the flow of multicast /broadcast MAC frames are kept within the associated VNET domain. All unicast MAC frames are sent across the standard LAN emulation virtual LAN. However, the interesting point is that each unicast frame will be addressed at layer 2 to stay within its VNET domain automatically. This means the edge devices are fast, inexpensive and simple while the VLAN/VNET architecture is optimized.

Detailed Description Text (8):
As represented by the cloud 105, a variety of other edge devices and ATM adapters may be coupled to the LAN emulation backbone 101, to establish a virtual LAN over a wide variety of LAN segments and across wide area links.

Detailed Description Text (15):
Standard LAN emulation ATM adapters can utilize this virtual net architecture. Because ATM adapters support a single end system, the configuration steps are not required for ATM end systems because there is no edge device connected to them which operates as a proxy for other LAN segments. Since the ATM adapters do not register as proxies within the LAN emulation server, the ATM adapter is automatically distinguished from an edge device. Thus, only one virtual channel per virtual LAN is set up from the server to the ATM end system. All associated virtual net membership entries set the virtual channel identifier values to the same virtual channel for the ATM adapter. In the configuration process, the virtual net membership entry, layer two MAC address is not forwarded to the ATM adapter, because it is not necessary for use there.

Detailed Description Text (24):
The bus which receives a multicast packet forwarded according to the process of FIG. 6, handles the packet according to the standards and distributes it to the edge devices within the virtual LAN.

Detailed Description Text (25):
The process shown in FIG. 7 is executed by edge devices receiving a packet from the BUS. Thus, the process begins with receiving a frame from the backbone network (block 310). The edge device next determines the type of frame (block 311). If it is a unicast frame, it is handled with the standard LAN emulation processes (block 312). If the frame is a multicast or a broadcast packet, the edge device determines whether the packet has a local source address (block 313). If it has a local source address, then it is discarded (block 314), because the process described in FIG. 6 has already forwarded the packet to the local LAN segments. If the packet does not have a local source address, then the edge device sends the frame once on ports to members of the virtual net membership list which matches the multicast packet (block 315). Accordingly, the virtual net server can be distributed to the edge devices in the emulated LAN in the ATM LAN emulation environment, or to similarly situated devices in other virtual LAN architectures.

h e b b g e e e f e e e g

Detailed Description Text (26):
Accordingly, the present invention provides management of traffic in a virtual LAN environment according to the concept of a virtual net domain. To maintain virtual net domain boundaries, edge devices operate at layer 2, while limited layer 3 complexity is centralized in, for example, an improved BUS server, or alternatively in a distributed virtual net server. Standard LAN emulation adapters are supported without modification. The virtual net server, the edge devices and adapters automatically learn virtual net domain membership among nodes on connected LAN segments.

Current US Original Classification (1):
709/223

CLAIMS:

1. For a network including a set of local area network (LAN) segments interconnected as a virtual LAN, in which nodes in the virtual LAN are members of one or more logical networks, a method for managing traffic in the network, comprising:

detecting a multi-destination packet on a LAN segment within the set;

determining, in response to the multi-destination packet, the logical network for which the detected multi-destination packet is intended; and

delivering information carried by the multi-destination packet to nodes authorized to receive multi-destination packets intended for members of the determined logical network.

38. A method for managing traffic in a virtual LAN environment in which there are a plurality of logical networks defined at layer three or above, comprising:

establishing virtual networks in the virtual LAN environment, virtual networks including nodes which are members of corresponding logical networks; and

confining, within the virtual LAN, multi-destination traffic originating from nodes within a given virtual network to nodes within the given virtual network.

41. The method of claim 38, wherein the step of establishing includes:

automatically learning the logical network of which nodes in the virtual LAN are members.

42. The method of claim 38, wherein the step of establishing includes:

automatically learning the logical network of which nodes in the virtual LAN are members, in response to multi-destination traffic transmitted by the nodes.

43. The method of claim 38, wherein the step of confining includes:

detecting a multi-destination packet within the virtual LAN;

determining the virtual network for which the multi-destination packet is intended; and

delivering the information in the multi-destination packet to the members of the determined virtual network by directed messages across the virtual LAN.